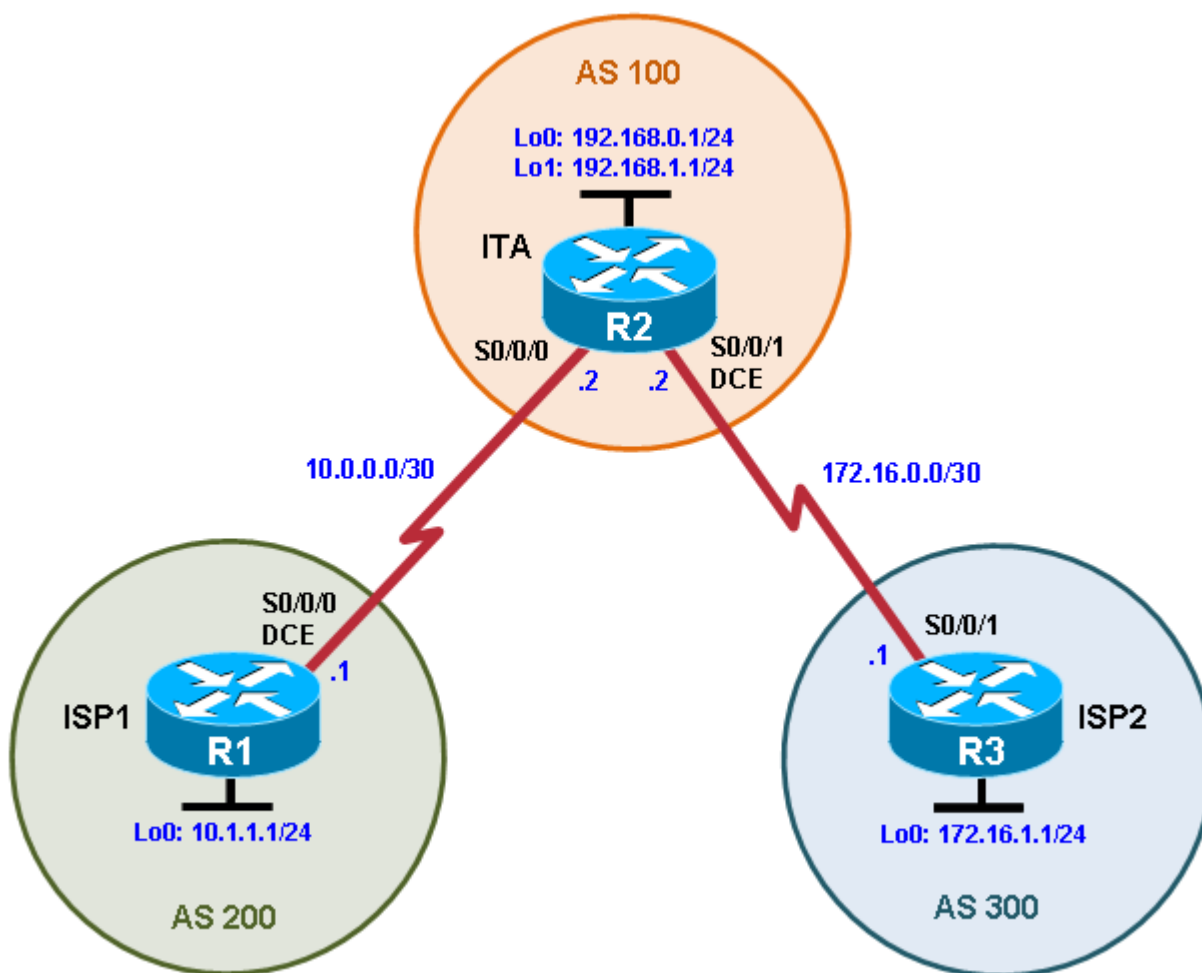


CCNPv7 ROUTE

Chapter 7 Lab 7-1, Configuring BGP with Default Routing

Topology



Objectives

- Configure BGP to exchange routing information with two ISPs.

Background

The International Travel Agency (ITA) relies extensively on the Internet for sales. For this reason, the ITA has decided to create a multihomed ISP connectivity solution and contracted with two ISPs for Internet connectivity with fault tolerance. Because the ITA is connecting to two different service providers, you must configure BGP, which runs between the ITA boundary router and the two ISP routers.

Note: This lab uses Cisco 1941 routers with Cisco IOS Release 15.4 with IP Base. The switches are Cisco WS-C2960-24TT-L with Fast Ethernet interfaces, therefore the router will use routing metrics associated with a 100 Mb/s

interface. Depending on the router or switch model and Cisco IOS Software version, the commands available and output produced might vary from what is shown in this lab.

Required Resources

- 3 routers (Cisco IOS Release 15.2 or comparable)
- Serial and Ethernet cables

Step 0: Suggested starting configurations.

- Apply the following configuration to each router along with the appropriate **hostname**. The **exec-timeout 0 0** command should only be used in a lab environment.

```
Router(config)# no ip domain-lookup
Router(config)# line con 0
Router(config-line)# logging synchronous
Router(config-line)# exec-timeout 0 0
```

Step 1: Configure interface addresses.

- Using the addressing scheme in the diagram, create the loopback interfaces and apply IPv4 addresses to these and the serial interfaces on ISP1 (R1), ISP2 (R3), and ITA (R2). The ISP loopbacks simulate real networks that can be reached through the ISP. The two loopbacks for the ITA router simulate the connections between the ITA boundary router and their core routers. Set a clock rate on the DCE serial interfaces.

```
ISP1(config)# interface Lo0
ISP1(config-if)# description ISP1 Internet Network
ISP1(config-if)# ip address 10.1.1.1 255.255.255.0
ISP1(config-if)# exit
ISP1(config)# interface Serial0/0/0
ISP1(config-if)# description ISP1 -> ITA
ISP1(config-if)# ip address 10.0.0.1 255.255.255.252
ISP1(config-if)# clock rate 128000
ISP1(config-if)# no shutdown
ISP1(config-if)# end
ISP1#

ITA(config)# interface Lo0
ITA(config-if)# description Core router network link 1
ITA(config-if)# ip address 192.168.0.1 255.255.255.0
ITA(config)# exit
ITA(config-if)# interface Lo1
ITA(config-if)# description Core router network link 2
ITA(config-if)# ip address 192.168.1.1 255.255.255.0
ITA(config-if)# exit
ITA(config)# interface Serial0/0/0
ITA(config-if)# description ITA -> ISP1
ITA(config-if)# ip address 10.0.0.2 255.255.255.252
ITA(config-if)# no shutdown
ITA(config-if)# exit
ITA(config)# interface Serial0/0/1
ITA(config-if)# description ITA -> ISP2
ITA(config-if)# ip address 172.16.0.2 255.255.255.252
```

```
ITA(config-if)# clock rate 128000
ITA(config-if)# no shutdown
ITA(config-if)# end
ITA#

ISP2(config)# interface Lo0
ISP2(config-if)# description ISP2 Internet Network
ISP2(config-if)# ip address 172.16.1.1 255.255.255.0
ISP2(config)# exit
ISP2(config-if)# interface Serial0/0/1
ISP2(config-if)# description ISP2 -> ITA
ISP2(config-if)# ip address 172.16.0.1 255.255.255.252
ISP2(config-if)# no shutdown
ISP2(config-if)# end
ISP2#
```

- b. Use **ping** to test the connectivity between the directly connected routers. Note that router ISP1 cannot reach router ISP2.

Step 2: Configure BGP on the ISP routers.

On the ISP1 and ISP2 routers, configure BGP to peer with the ITA boundary router and advertise the ISP loopback networks.

```
ISP1(config)# router bgp 200
ISP1(config-router)# neighbor 10.0.0.2 remote-as 100
ISP1(config-router)# network 10.1.1.0 mask 255.255.255.0

ISP2(config)# router bgp 300
ISP2(config-router)# neighbor 172.16.0.2 remote-as 100
ISP2(config-router)# network 172.16.1.0 mask 255.255.255.0
```

Step 3: Configure BGP on the ITA boundary router.

- a. Configure the ITA router to run BGP with both Internet providers.

```
ITA(config)# router bgp 100
ITA(config-router)# neighbor 10.0.0.1 remote-as 200
ITA(config-router)# neighbor 172.16.0.1 remote-as 300
ITA(config-router)# network 192.168.0.0
ITA(config-router)# network 192.168.1.0
```

You should see BGP neighbor peering messages on the console similar to the following.

```
*Sep  8 16:00:21.587: %BGP-5-ADJCHANGE: neighbor 10.0.0.1 Up
```

- b. To verify the configuration, check the ITA routing table with the **show ip route** command.

```
ITA# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override
```

Gateway of last resort is not set

```

      10.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
C       10.0.0.0/30 is directly connected, Serial0/0/0
L       10.0.0.2/32 is directly connected, Serial0/0/0
B       10.1.1.0/24 [20/0] via 10.0.0.1, 00:01:10
      172.16.0.0/16 is variably subnetted, 3 subnets, 3 masks
C       172.16.0.0/30 is directly connected, Serial0/0/1
L       172.16.0.2/32 is directly connected, Serial0/0/1
B       172.16.1.0/24 [20/0] via 172.16.0.1, 00:00:53
      192.168.0.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.0.0/24 is directly connected, Loopback0
L       192.168.0.1/32 is directly connected, Loopback0
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, Loopback1
L       192.168.1.1/32 is directly connected, Loopback1
ITA#

```

ITA has BGP routes to the loopback networks at each ISP router.

- c. Run the following Tcl script on all routers to verify connectivity. If these pings are not successful, troubleshoot. Use **exit** to exit the Tcl script.

Note: The WAN subnets connecting ITA (R2) to the ISPs (R1 and R3) are not advertised in BGP, so the ISPs will not be able to ping each other's serial interface address.

```

ITA# tclsh

foreach address {
10.0.0.1
10.0.0.2
10.1.1.1
172.16.0.1
172.16.0.2
172.16.1.1
192.168.0.1
192.168.1.1
} {
ping $address }

```

Step 4: Verify BGP on the routers.

- a. To verify the BGP operation on ITA, issue the **show ip bgp** command.

```

ITA# show ip bgp
BGP table version is 5, local router ID is 192.168.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop           Metric LocPrf Weight Path
*> 10.1.1.0/24      10.0.0.1             0                0 200 i
*> 172.16.1.0/24    172.16.0.1           0                0 300 i
*> 192.168.0.0      0.0.0.0              0               32768 i
*> 192.168.1.0      0.0.0.0              0               32768 i
ITA#

```

What is the local router ID?

Which table version is displayed?

An asterisk (*) next to a route indicates that it is valid. An angle bracket (>) indicates that the route has been selected as the best route.

- b. To verify the operation of ISP1, issue the **show ip bgp** command.

```
ISP1# show ip bgp
BGP table version is 5, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	10.1.1.0/24	0.0.0.0	0		32768	i
*>	172.16.1.0/24	10.0.0.2			0 100 300	i
*>	192.168.0.0	10.0.0.2	0		0 100	i
*>	192.168.1.0	10.0.0.2	0		0 100	i

ISP1#

Which table version is displayed and is it the same as the BGP table version for ITA?

From ISP1, what is the path to network 172.16.1.0/24?

- c. On the ISP1 router, issue the **shutdown** command on Loopback0. Then on ITA, issue the **show ip bgp** command again.

```
ISP1(config)# interface loopback 0
ISP1(config-if)# shutdown
ISP1(config-if)#
```

```
ITA# show ip bgp
BGP table version is 6, local router ID is 192.168.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	172.16.1.0/24	172.16.0.1	0		0 300	i
*>	192.168.0.0	0.0.0.0	0		32768	i
*>	192.168.1.0	0.0.0.0	0		32768	i

ITA#

Which table version is displayed? Why?

What happened to the route for network 10.1.1.0/24?

- d. Bring ISP1 router Loopback0 back up by issuing the **no shutdown** command.

```
ISP1(config)# interface loopback 0
ISP1(config-if)# no shutdown
ISP1(config-if)#
```

- e. On ITA, issue the **show ip bgp neighbors** command. The following is a partial sample output of the command showing neighbor 172.16.0.1.

```
ITA# show ip bgp neighbors
BGP neighbor is 10.0.0.1, remote AS 200, external link
  BGP version 4, remote router ID 10.1.1.1
  BGP state = Established, up for 00:20:47
  Last read 00:00:49, last write 00:00:41, hold time is 180, keepalive interval is
60 seconds
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Enhanced Refresh Capability: advertised and received
    Multisession Capability:
    Stateful switchover support enabled: NO for session 1
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

              Sent          Rcvd
Opens:              1           1
Notifications:      0           0
Updates:             5           1
Keepalives:         15          17
Route Refresh:       0           0
Total:              21          19
Default minimum time between advertisement runs is 30 seconds
```

<output omitted>

Based on the output of this command, what is the BGP state between this router and ISP2?

How long has this connection been up?

Step 5: Configure route filters.

- a. Check the ISP2 routing table using the **show ip route** command. ISP2 should have a route that belongs to ISP1, network 10.1.1.0.

```
ISP2# show ip route
<output omitted>
```

```

10.0.0.0/24 is subnetted, 1 subnets
B    10.1.1.0 [20/0] via 172.16.0.2, 00:09:26
172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
C    172.16.0.0/30 is directly connected, Serial0/0/1
L    172.16.0.1/32 is directly connected, Serial0/0/1
C    172.16.1.0/24 is directly connected, Loopback0
L    172.16.1.1/32 is directly connected, Loopback0
B    192.168.0.0/24 [20/0] via 172.16.0.2, 00:28:05
B    192.168.1.0/24 [20/0] via 172.16.0.2, 00:28:05
ISP2#

```

If ITA advertises a route belonging to ISP1, ISP2 installs that route in its table. ISP2 might then attempt to route transit traffic through the ITA. This would make ITA a transit router. A traceroute to ISP1's Lo0 interface illustrates this issue.

```

ISP2# traceroute 10.1.1.1
Type escape sequence to abort.
Tracing the route to 10.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
 1 172.16.0.2 8 msec 4 msec 8 msec
 2 * * *
 3 * * *
 4 * * * <control-shift-6 to break>
ISP2#

```

The **traceroute 10.1.1.1** fails because ISP1 does not have a route to the source IPv4 address of the traceroute, 172.16.0.1. It is common in BGP networks not to advertise the links between providers in BGP. A traceroute using the source IPv4 address of ISP2' Lo0 interface is successful, showing that ITA is a transit router for this network.

```

ISP2# traceroute 10.1.1.1 source loopback0
Type escape sequence to abort.
Tracing the route to 10.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
 1 172.16.0.2 8 msec 4 msec 8 msec
 2 10.0.0.1 12 msec * 12 msec
ISP2#

```

- b. Configure the ITA router so that it advertises only ITA networks 192.168.0.0 and 192.168.1.0 to both providers. On the ITA router, configure the following access list.

```
ITA(config)# access-list 1 permit 192.168.0.0 0.0.1.255
```

- c. Apply this access list as a route filter using the **distribute-list** keyword with the BGP **neighbor** statement.

```

ITA(config)# router bgp 100
ITA(config-router)# neighbor 10.0.0.1 distribute-list 1 out
ITA(config-router)# neighbor 172.16.0.1 distribute-list 1 out

```

- d. Check the routing table for ISP2 again. The route to 10.1.1.0, ISP1, should still be in the table.

```

ISP2# show ip route
<output omitted>

```

```

10.0.0.0/24 is subnetted, 1 subnets
B    10.1.1.0 [20/0] via 172.16.0.2, 00:25:14
172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
C    172.16.0.0/30 is directly connected, Serial0/0/1

```

```

L      172.16.0.1/32 is directly connected, Serial0/0/1
C      172.16.1.0/24 is directly connected, Loopback0
L      172.16.1.1/32 is directly connected, Loopback0
B      192.168.0.0/24 [20/0] via 172.16.0.2, 00:43:53
B      192.168.1.0/24 [20/0] via 172.16.0.2, 00:43:53
ISP2#

```

- e. Return to ITA and issue the **clear ip bgp *** command. Wait until the routers reach the established state, which might take several seconds, and then recheck the ISP2 routing table. The route to ISP1, network 10.1.1.0, should no longer be in the routing table for ISP2, and the route to ISP2, network 172.16.1.0, should not be in the routing table for ISP1.

```

ITA# clear ip bgp *
ITA#
*Sep  8 16:47:25.179: %BGP-5-ADJCHANGE: neighbor 10.0.0.1 Down User reset
*Sep  8 16:47:25.179: %BGP_SESSION-5-ADJCHANGE: neighbor 10.0.0.1 IPv4 Unicast
topology base removed from session User reset
*Sep  8 16:47:25.179: %BGP-5-ADJCHANGE: neighbor 172.16.0.1 Down User reset
*Sep  8 16:47:25.179: %BGP_SESSION-5-ADJCHANGE: neighbor 172.16.0.1 IPv4 Unicast
topology base removed from session User reset
*Sep  8 16:47:25.815: %BGP-5-ADJCHANGE: neighbor 10.0.0.1 Up
*Sep  8 16:47:25.819: %BGP-5-ADJCHANGE
ITA#: neighbor 172.16.0.1 Up
ITA#

```

Note: The **clear ip bgp *** command is disruptive because it completely resets all BGP adjacencies. This is acceptable in a lab environment but could be problematic in a production network. Instead, if only a change of inbound/outbound routing policies is to be performed, it is sufficient to issue the **clear ip bgp * in** or **clear ip bgp * out** commands. These commands perform only a new BGP database synchronization without the disruptive effects of a complete BGP adjacency reset. All current Cisco IOS versions support the route refresh capability that replaces the inbound soft reconfiguration feature that previously had to be configured on a per-neighbor basis.

```

ISP2# show ip route
<output omitted>

```

```

      172.16.0.0/16 is variably subnetted, 4 subnets, 3 masks
C      172.16.0.0/30 is directly connected, Serial0/0/1
L      172.16.0.1/32 is directly connected, Serial0/0/1
C      172.16.1.0/24 is directly connected, Loopback0
L      172.16.1.1/32 is directly connected, Loopback0
B      192.168.0.0/24 [20/0] via 172.16.0.2, 00:00:06
B      192.168.1.0/24 [20/0] via 172.16.0.2, 00:00:06
ISP2#

```

```

ISP1# show ip route
<output omitted>

```

```

      10.0.0.0/8 is variably subnetted, 4 subnets, 3 masks
C      10.0.0.0/30 is directly connected, Serial0/0/0
L      10.0.0.1/32 is directly connected, Serial0/0/0
C      10.1.1.0/24 is directly connected, Loopback0
L      10.1.1.1/32 is directly connected, Loopback0
B      192.168.0.0/24 [20/0] via 10.0.0.2, 00:00:42
B      192.168.1.0/24 [20/0] via 10.0.0.2, 00:00:42

```



```
ISP1#
```

Step 6: Configure primary and backup routes using floating static routes.

With bidirectional communication established with each ISP via BGP, configure the primary and backup routes. This can be done with floating static routes or BGP.

- Issue the **show ip route** command on the ITA router.

```
ITA# show ip route
<output omitted>
```

```
Gateway of last resort is not set
```

```

      10.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
C       10.0.0.0/30 is directly connected, Serial0/0/0
L       10.0.0.2/32 is directly connected, Serial0/0/0
B       10.1.1.0/24 [20/0] via 10.0.0.1, 00:03:51
      172.16.0.0/16 is variably subnetted, 3 subnets, 3 masks
C       172.16.0.0/30 is directly connected, Serial0/0/1
L       172.16.0.2/32 is directly connected, Serial0/0/1
B       172.16.1.0/24 [20/0] via 172.16.0.1, 00:03:51
      192.168.0.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.0.0/24 is directly connected, Loopback0
L       192.168.0.1/32 is directly connected, Loopback0
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, Loopback1
L       192.168.1.1/32 is directly connected, Loopback1
ITA#
```

Notice that there is no gateway of last resort defined. This is a problem because ITA is the border router for the corporate network.

- Configure static routes to reflect the policy that ISP1 is the primary provider and that ISP2 acts as the backup by specifying a lower distance metric for the route to ISP1 (210) as compared to the backup route to ISP2 (distance metric 220).

```
ITA(config)# ip route 0.0.0.0 0.0.0.0 10.0.0.1 210
ITA(config)# ip route 0.0.0.0 0.0.0.0 172.16.0.1 220
```

- Verify that a default route is defined using the **show ip route** command.

```
ITA# show ip route
<output omitted>
Gateway of last resort is 10.0.0.1 to network 0.0.0.0
```

```

S*      0.0.0.0/0 [210/0] via 10.0.0.1
      10.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
C       10.0.0.0/30 is directly connected, Serial0/0/0
L       10.0.0.2/32 is directly connected, Serial0/0/0
B       10.1.1.0/24 [20/0] via 10.0.0.1, 00:05:38
      172.16.0.0/16 is variably subnetted, 3 subnets, 3 masks
C       172.16.0.0/30 is directly connected, Serial0/0/1
L       172.16.0.2/32 is directly connected, Serial0/0/1
B       172.16.1.0/24 [20/0] via 172.16.0.1, 00:05:38
      192.168.0.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.0.0/24 is directly connected, Loopback0
L       192.168.0.1/32 is directly connected, Loopback0
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
```

```

C      192.168.1.0/24 is directly connected, Loopback1
L      192.168.1.1/32 is directly connected, Loopback1
ITA#

```

- d. Test this default route by creating an unadvertised loopback on the router for ISP1.

```

ISP1# config t
ISP1(config)# interface loopback 100
ISP1(config-if)# ip address 192.168.100.1 255.255.255.0

```

- e. Issue the **show ip route** command to ensure that the newly added 192.168.100.0/24 network does not appear in the routing table.

```

ITA# show ip route
<output omitted>

```

Gateway of last resort is 10.0.0.1 to network 0.0.0.0

```

S*    0.0.0.0/0 [210/0] via 10.0.0.1
      10.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
C      10.0.0.0/30 is directly connected, Serial0/0/0
L      10.0.0.2/32 is directly connected, Serial0/0/0
B      10.1.1.0/24 [20/0] via 10.0.0.1, 00:07:08
      172.16.0.0/16 is variably subnetted, 3 subnets, 3 masks
C      172.16.0.0/30 is directly connected, Serial0/0/1
L      172.16.0.2/32 is directly connected, Serial0/0/1
B      172.16.1.0/24 [20/0] via 172.16.0.1, 00:07:08
      192.168.0.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.0.0/24 is directly connected, Loopback0
L      192.168.0.1/32 is directly connected, Loopback0
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.1.0/24 is directly connected, Loopback1
L      192.168.1.1/32 is directly connected, Loopback1
ITA#

```

- f. In extended ping mode, ping the ISP1 loopback 1 interface 192.168.100.1 with the source originating from the ITA loopback 1 interface 192.168.1.1.

```

ITA# ping
Protocol [ip]:
Target IP address: 192.168.100.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 192.168.1.1
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.100.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.1.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/14/16 ms
ITA#

```

Note: You can bypass extended ping prompted mode and ping while specifying a source address using one of these abbreviated commands:

```
ITA# ping 192.168.100.1 source 192.168.1.1
```

or

```
ITA# ping 192.168.100.1 source Lo1
```

Note: Testing the default route by creating an unadvertised network on ISP1 and pinging it works only because the default route also points toward ISP1. If the preferred default route pointed toward ISP2, the ping to that unadvertised network on ISP1 would not succeed. If the link to ISP1 failed, the default route to ISP2 would become active, but the pings would be successful only if ISP1 and ISP2 have another working interconnection and appropriate BGP peering between them, which is currently not the case.

Step 7: Using BGP to propagate a default route.

- a. ISP router will be used to inject a default route via BGP. First, remove the current default routes on ITA.

```
ITA(config)# no ip route 0.0.0.0 0.0.0.0 10.0.0.1 210
ITA(config)# no ip route 0.0.0.0 0.0.0.0 172.16.0.1 220
```

- b. Next, configure the ISP1 router to send a default route to its neighbor, the ITA router. This command does not require the presence of 0.0.0.0 in the local ISP1 router.

```
ISP1(config)# router bgp 200
ISP1(config-router)# neighbor 10.0.0.2 default-originate
ISP1(config-router)#
```

- c. Verify that the default route was received by ITA using BGP.

```
ITA# show ip route
<output omitted>
```

```
Gateway of last resort is 10.0.0.1 to network 0.0.0.0
```

```
B* 0.0.0.0/0 [20/0] via 10.0.0.1, 00:01:43
    10.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
C    10.0.0.0/30 is directly connected, Serial0/0/0
L    10.0.0.2/32 is directly connected, Serial0/0/0
B    10.1.1.0/24 [20/0] via 10.0.0.1, 00:06:51
    172.16.0.0/16 is variably subnetted, 3 subnets, 3 masks
C    172.16.0.0/30 is directly connected, Serial0/0/1
L    172.16.0.2/32 is directly connected, Serial0/0/1
B    172.16.1.0/24 [20/0] via 172.16.0.1, 00:06:51
    192.168.0.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.0.0/24 is directly connected, Loopback0
L    192.168.0.1/32 is directly connected, Loopback0
    192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.1.0/24 is directly connected, Loopback1
L    192.168.1.1/32 is directly connected, Loopback1
ITA#
```